Optimizing Component Combination in a Multi-Indexing Paragraph Retrieval System

Boris Iolis
Université Libre de Bruxelles (ULB)
Boulevard du Triomphe
Campus de la Plaine, CP212
B-1050 Brussels - Belgium
biolis@ulb.ac.be

Gianluca Bontempi
Université Libre de Bruxelles (ULB)
Boulevard du Triomphe
Campus de la Plaine, CP212
B-1050 Brussels - Belgium
gbonte@ulb.ac.be

arXiv:1408.2430v1 [cs.IR] 11 Aug 2014

Abstract

We demonstrate a method to optimize the combination of distinct components in a paragraph retrieval system. Our system makes use of several indices, query generators and filters, each of them potentially contributing to the quality of the returned list of results. The components are combined with a weighed sum, and we optimize the weights using a heuristic optimization algorithm. This allows us to maximize the quality of our results, but also to determine which components are most valuable in our system. We evaluate our approach on the paragraph selection task of a Question Answering dataset.

1 Introduction

In the Information Retrieval domain, the combination of search results is a long studied problem, and can effectively increase the precision of the resulting system. As a result, many IR systems are designed to use multiple querying methods, and then combine the retrieved results, which is also called Data Fusion. Fox and Shaw (1994) showed the effectiveness of combining multiple retrieval runs as opposed to selecting only one of them. Lee (1995) combined search strategies using a simple, non-weighed sum. Vogt and Cotrell (1999) used a Linear Combination model, for which they optimized the weights to maximize the system's precision. Tiedemann (2005) employs a Genetic Algorithm to perform a similar optimization.

In this paper, we demonstrate a method to optimally combine the components of a paragraph retrieval system. Our approach is similar to those listed above, as we use a simple Linear Combination model: however, we include all our system's components in this model, and not only the querying modules. This allows us to not only optimize the mixture of querying methods, but also of filters and scorers.

We built our system using a multi-indexing architecture (several indices are being used for the same text corpus), and including some state-of-the-art query generators and filters. Our system retrieves a set of paragraphs from the text corpus based on the input query, which are then ranked according to their confidence scores and constitute the results list. All components are treated on the same level, and can equally contribute to the final confidence score associated to each retrieved paragraph. The components are combined using the Linear Combination model, and its weights are tuned using a Heuristic Optimization algorithm. Finally, we evaluate the results on a paragraph selection task using a Question Answering dataset.

The rest of the paper is organized as follows: in the next section, we present the architecture of our paragraph retrieval system. In Section 3, we explain how all the components can be combined and tuned. Finally, Section 4 presents our experimental results, while Section 5 contains the conclusion and discussion on future work.

2 System Architecture

The architecture of our paragraph retrieval system is illustrated in Figure 1. It is based on the typical design of a Question Answering system (see for instance (Hirschman and Gaizauskas, 2002)), without answer extraction, as we try to retrieve a paragraph containing the correct answer to an input question instead of extracting the exact answer string from the text. We do however use multi-indexing, which is, to our knowledge, not so commonly studied in QA literature.

Starting from a single text corpus, we create a set of indices which will be used for querying. For each index, the text corpus is pre-processed in a distinct way. So far, our system implements the following four indices:

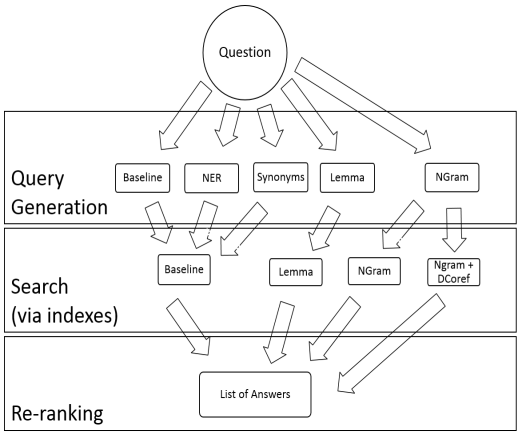

Figure 1: Our multi-indexing search system architecture.

- **Baseline**: standard inverted index on the text corpus, built using Lucene, which includes stopwords removal and simple tokenization. Only unigrams are indexed in this case
- **Lemmatization**: same as the baseline index but with a lemmatization step applied to the text corpus during pre-processing
- **NGrams**: same as the baseline index but with 2-grams and 3-grams added to the indexing terms
- **NGrams + Coreference Resolution**: same as NGrams, but with a coreference resolution step before indexing

We chose this multi-indexing approach in order to maximize the probability of retrieving the right paragraph in the querying stage (through at least one index). Typically, indexing in any specific way has its pros and cons; either we generalize too much (linking many similar terms to the same indexed term, for instance their common lemma), or not enough (indexing all words or ngrams separately). As queries will sometimes work better with more generalization, and sometimes with less, we are trying to get the best of both worlds by creating multiple indices and using them in parallel. Although there is a cost associated to creating and maintaining multiple indices, both in terms of disk space and pre-processing time, we believe that, even if the resulting improvements in recall are minimal, the benefits will outweigh the costs as long as the number of indices used is not excessively large.

In the querying stage, the input question is transformed into several queries, which is a common technique in IR and QA (see for instance (Dumais et al., 2002)). One query is generated for each index to match its specificities. For instance, to query the lemmatization index, the input question needs to undergo the same lemmatization step as did the text corpus. Furthermore, two additional query generation approaches are implemented, and both are used on the baseline index;

- **Named Entity Recognition**: builds a query containing only the named entities found in the input question
- **Synonyms**: query expansion with synonyms based on WordNet (Miller, 1995)

Each query will return a list of paragraphs; in the last stage of our system, those paragraphs will be evaluated using a set of criteria, and then re-ranked in order to provide the most relevant list of paragraphs with regards to the original question. This re-ranking is based on our scoring framework, which is presented in the next section. The criteria we use at this stage are based on word counts (used extensively in IR and QA literature, for instance in (Ligozat et al., 2006)) and Latent Dirichlet Allocation (LDA) (Blei and Jordan, 2003).

- **Common words**: number of common words between the paragraph and the input question
- **Common 2-grams**: same as above but with 2-grams instead of single words
- **Common 3-grams**: same as above but with 3-grams
- **LDA-10**: cosine similarity between the probability vectors of the paragraph and the input question, based on a LDA model with 10 topics, trained on the text corpus
- **LDA-100**: same as above but with a 100 topics model

3 Scoring Framework

3.1 Computing the Scores

Our method to score a paragraph is a simple application of the Linear Combination model to all the components of our system. We compile a list of criteria (we will call them features through the rest

of the paper) consisting of all the query generators from the querying stage and the evaluators from the re-ranking stage. Each of those features gives a distinct score to each paragraph. For queries, the score of a paragraph is given by the Lucene confidence score if this paragraph was returned in the results list when using this query, and it is set to 0 otherwise. For evaluators, this is straightforward. Each of these scores is then normalized using the Z-score normalization method (Montague and Aslam, 2001). Finally, the overall score of a paragraph c_i is computed as a linear combination of the features f_j, as shown below:

$$Score_{c_i} = \sum_{j=1}^{N} w_j \times f_j(c_i) \qquad (1)$$

where N is the number of components (evaluators and query generators) of the system (in our case $N = 11$); $f_j(c_i)$ is the score given by component j to paragraph i; and w_j are weights such that $\sum_{j=1}^{N} w_j = 1$

The actual ranking of the paragraph can be done by simply sorting them according to their score. This approach allows us to easily combine all the components of our system to obtain a global score for each paragraph.

3.2 Tuning the Weights

In (1), the weights should be tuned to maximize precision. They could be defined manually according to the quality of each feature (how relevant are the scores given by the feature), but unfortunately we do not have this knowledge beforehand. Also, evaluating each feature individually does not account for their diversity and complementarity when combined. Therefore, we decided to treat the tuning of those weights as a multivariate optimization problem, where the objective is to find the set of weights w_j maximizing the overall performance of the system, according to an evaluation metric of interest. Though the cost function is not differentiable, we can still apply a wide variety of heuristic optimization methods (coordinate ascent, simulated annealing,...) to find the (approximate) best set of weights. For this work, we used a Differential Evolution algorithm (Storn and Price, 1997) to perform this task, as it would allow us to demonstrate the effectiveness of our approach while being relatively simple to implement.

4 Experimental Results

4.1 Dataset

We used the dataset from the ResPubliQA 2010 competition (Penas et al., 2010), containing a text corpus of 10,700 European parliament transcripts (taken from the JRC-Acquis[1] and Europarl[2] collections), accompanied with a set of 200 questions, each having the correct answer provided (gold standard). The text documents are structured in numbered paragraphs of a few sentences each. We focused on the paragraph selection task (finding the paragraph containing the correct answer), which made it possible to perform automated assessment, by comparing the identifiers of the retrieved paragraphs to the gold standard. We compared our results with the work of (Molino et al., 2012), who perform the same paragraph selection task on the same dataset.

4.2 Results

Table 1 shows the results obtained by our system, first with all components combined in a naive way (all weights w_j from (1) being equal), and then with weight tuning as described in Section 3.2. Our metric of choice is the Mean Reciprocal Rank (MRR), which gives a score of $1/r$ for each question, where r is the position of the paragraph containing the right answer in the results list. For the weight tuning experiment, we used 20 rounds of cross-validation to avoid over-fitting. In each round, the tuning was done on 190 questions, and then evaluated on the remaining 10. The result shown in the table is the average of those 20 MRR scores.

System	MRR
Our system	0.513
Our system (with weight tuning)	0.543
QuestionCube (baseline)	0.549
QuestionCube (best)	0.637

Table 1: Evaluation of our system, with and without weight tuning, and comparison with the QuestionCube system from (Molino et al., 2012).

We see that our system performs better when the combination of components is tuned with the Differential Evolution algorithm. Furthermore, our

[1] http://ipsc.jrc.ec.europa.eu/index.php?id=198
[2] http://www.statmt.org/europarl/

Component	Average Weight
Baseline	0.0
Lemmatization	0.303
NGrams	0.0
NGram + Coref.	0.033
Named Entities	0.0
Synonyms	0.357
Common Unigrams	0.0
Common 2-Grams	0.253
Common 3-Grams	0.054
LDA-10	0.0
LDA-100	0.0

Table 2: Average weights given to each component across the 20 cross-validation runs of the Differential Evolution algorithm

results are not so far behind the QuestionCube system. We are in line with the performance of their baseline version (which is already a full-fledged QA system on its own), but are behind the improved version from (Molino et al., 2012), which uses far more advanced distributional semantic models than our simple LDA evaluators.

Finally, our parameter tuning experiment gives us some insights on the added value of each component in our system; if a component is consistently given a weight of 0 by the optimization algorithm, we can conclude it is not very valuable for the overall performance of our system. The average weight for each component across the 20 cross-validation runs are shown in Table 2. We can see that more than 90% of the total weight was concentrated among three specific features: Lemmatization, Synonyms, and the 2-Grams evaluator. This preference for a very limited subset of components might suggest that our choice of components to implement might not have been the best, or that some of them might require additional fine-tuning.

5 Conclusion

In this paper, we demonstrated our method to efficiently combine the components of a paragraph retrieval system. We showed that using a heuristic optimization algorithm to tune this combination had a positive effect on the performance of our system. The overall performance is also in line with previous evaluations on the same dataset. Finally, we showed how this methodology could be used to evaluate the added value of each component which could be useful in our future work.

Now that we have this framework as a backbone, we can easily add new components to the system to make it more competitive in the future, as only the basic components have been integrated so far. As was shown in the weight tuning experiment, some effort may be required to understand why some of our components do not bring so much added value, and modify them to address this situation. Different optimization methods could also be implemented.

Acknowledgments

We would like to thank David Verborgh for providing valuable input during the development of our application. This work is part of a PhD project funded by the Innoviris institute, via their Doctiris program, and carried out in cooperation with Mentis[3].

References

Blei, D., Ng, A., Jordan, M.: Latent Dirichlet Allocation. The Journal of Machine Learning Research, 3, p.993-1022, (2003)

Chu-Carroll, J., Czuba, K., Prager, J., Itterycheriah, A.: In Question Answering , Two Heads Are Better Than One. NAACL 03 Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, Volume 1, 2431 (2003)

Croft, B.: Combining Approaches to Information Retrieval. Advances in Information Retrieval, 136 (2002)

Dumais, S., Banko, M., Brill, E., Lin, J., Ng, A.: Web Question Answering : Is More Always Better? Proceedings of the 25th annual international ACM SIGIR conference on research and development in information retrieval (2002)

Fox, E.A., Shaw, J.A.: Combination of Multiple Searches. Proceedings of the Second Text REtrieval Conference (TREC-2) (1994)

Hirschman, L., Gaizauskas, R.: Natural Language Question Answering: The View from Here. Natural Language Engineering, 7(04), 275300 (2002)

Lee, J.: Combining Multiple Evidence from Different Properties of Weighting Schemes. SIGIR '95 Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval (1995)

[3]http://www.mentis-consulting.be/

Ligozat, A., Grau, B., Vilnat, A., Robba, I., Grappy, A.: Lexical Validation of Answers in Question Answering. IEEE/WIC/ACM International Conference on Web Intelligence (2006)

Miller, G. A.: Wordnet: A Lexical Database for English. ACM (1995)

Molino, P., Basile, P., Caputo, A., Lops, P., Semeraro, G.: Exploiting Distributional Semantic Models in Question Answering. IEEE Sixth International Conference on Semantic Computing, 146153 (2012)

Montague, M., Aslam, J. A.: Relevance Score Normalization for Metasearch. Proceedings of the ACM CIKM, 427433 (2001)

Penas, A., Forner, P., Rodrigo, A., Sutcliffe, R., Forascu, C., Mota, C.: Overview of ResPubliQA 2010 : Question Answering Evaluation over European Legislation. Working notes of ResPubliQA 2010, Conference and Labs of the Evaluation Forum (CLEF) (2010)

Storn, R., Price, K.: Differential Evolution A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. Journal of Global Optimization, Volume 11 Issue 4, 341-359 (1997)

Tiedemann, J.: Improving Passage Retrieval in Question Answering using NLP. Proceedings of the 12th Portuguese Conference on Artificial Intelligence (EPIA) (2005)

Vogt, C., Cottrell, G.: Fusion Via a Linear Combination of Scores. Information Retrieval, October 1999, Volume 1, Issue 3, pp 151-173 (1999)